

SQL Syntax & Query Performance

By Joe Celko

For Embarcadero Technologies
October, 2015

Table of Contents

SQL Syntax & Query Performance	3
Write SQL That Looks Like SQL.....	3
BETWEEN Predicates	3
Membership Predicates	4
Generalized Group Predicates	5
The <theta> ALL <subquery> Syntax	6
The ALL Predicate and Extrema Functions	7
EXISTS().....	8
The OVERLAPS() Predicate.....	11
The CASE Expression	12
Do Not Retrieve More Columns Than You Need.....	13
DELETE and UPDATE in Batches	14
Use Stored Procedures	14
GUIDs are Not What You Think.....	15
Do Not Use Count() if You Need Existence	15
Do Not Do Negative Searches.....	16
Use CTEs	16
“Double-Dipping” and the Windows Clause.....	17
Conclusions	19
About The Author.....	20

SQL Syntax & Query Performance

The 2000 book *Me Talk Pretty One Day* (ISBN 978-0316777728) is a collection of essays by American humorist David Sedaris. He moved to a rural part of Normandy, France, and had no idea how to speak or read French at the time. The humor comes from his frustrated attempts to learn the language and culture. The locals knew him as the American idiot who scares their animals by screaming. The warning on electrified fences was in French and a great source of amusement for the neighbors.

New SQL developers on every platform also do not “talk pretty” and keep struggling trying get out of the mindset of their prior language. The SQL language made a huge jump in the SQL-92 Standards, adding many new language features based on the actual usage of the language, improvements in the hardware and attempts to make “talk SQL pretty” easier and more natural.

The query that runs well on a local set of test data can fail miserably on the production system. Very often a simple change in the syntax can make a huge performance difference. Much of SQL advice is “it depends” because tuning a database is both art and science.

However, there are some principles you can follow that should produce good SQL most of the time and give you a maintainable schema. These are heuristics, not mathematical formulas.

Write SQL That Looks Like SQL

FORTRAN, BASIC, COBOL and most other languages have a limited set of logical operators based on simple Boolean (i.e. two valued) algebra. The problem is not just the SQL has a three valued logic, but it has short hands that deal with “sets of predicates” instead of a linear sequence of infix Boolean expressions.

BETWEEN Predicates

Consider this simple generic logical expression, which occurs in real code all the time.

```
((a <= x) and (x >= b))
```

This will work just fine in SQL, too. *But an SQL programmer would use:*

```
x BETWEEN a AND b
```

See the difference in the mindsets? A pair of two very simple predicates is reduced to a 3-ary operator. Between-ness started as a geometric property for points on a line. The optimizations that use it are based on those theorems.

Now that the higher level concept is locked in a bit of syntax, we can generalize it with extra syntax. The BNF for the BETWEEN is as follows, and the asymmetric option is the default equivalent to the code I just gave.

```
<between predicate>  
::= <row value predicand> <between predicate part 2>
```

```
<between predicate part 2> ::=  
[NOT] BETWEEN [ASYMMETRIC, SYMMETRIC]  
<row value predicand> AND <row value predicand>
```

```
X BETWEEN SYMMETRIC y AND z
```

means

```
((x BETWEEN ASYMMETRIC y AND z)  
OR (x BETWEEN ASYMMETRIC z AND y))
```

See how we are getting a lot of programming power with minor syntax extensions as we expand the language? Try another example of this principle:

Membership Predicates

Consider this simple generic logical expression, which you will see in a lot of programs:

```
((foo = 1) or (foo = 2) or (foo = 3))
```

This will work just fine in SQL, too. *But an SQL programmer would use:*

```
foo IN (1,2,3)
```

We stole this syntax for the simple IN() from Pascal. It is an n-ary ($n > 0$) logical operator. Is there any advantage to using this over a chain of OR-ed expressions? The first obvious advantage is the amount of typing you can save. The next advantage is that the

right-hand list is much easier to read; you see that foo is compared to an element in a *set of values*. This is the conceptual difference between a bunch of trees and a forest, employees versus the abstraction of Personnel, what Cantor called "the many that is treated as a one" when he defined set theory, and so forth.

The final advantage is that some SQL product will *optimize the search of the list*. The values can be put into a binary tree or a hash table to speed up searching when the list is long enough. This leads to the counter-intuitive situation where adding extra elements to the list makes the query run faster by making the list long enough for the optimization to kick in. These extra values can be absurd, so the code looks really strange.

This syntax is easy to generalize to any expression that returns a list or a one-column table like "x IN (SELECT y FROM ..)" But there is other syntax for this that most programmers do not know: It generalizes the theta operator rather than the source of the list.

Generalized Group Predicates

A theta operators is the fancy name for a comparison operator and can be any of { =, <, >, <>, <=, >= } and not just the equality of the IN() predicate. In fact, the IN() predicate is defined in the ANSI/ISO Standard as a short hand for "= ANY" predicate and "<> ANY" as a short hand for NOT IN().

There are two shorthand notations that compare a single value to those of a subquery, and take the general form <value expression> <theta op> <quantifier> <subquery> . The first predicate "<value expression> <theta op> [ANY, SOME] <table expression> " is equivalent to taking each row, s, (assume that they are numbered from 1 to n) of <table expression> and testing "<value expression> <theta op> s" with ORs between the expanded expressions:

```
((<value expression> <theta op> s1)
OR (<value expression> <theta op> s2)
...
OR (<value expression> <theta op> sn))
```

When you get a single TRUE result, the whole predicate is TRUE. As long as <table expression> has cardinality greater than zero and one non-NULL value, you will get a result of TRUE or FALSE. The keyword SOME is the same as ANY, and the choice is just a matter of style and readability.

The <theta> ALL <subquery> Syntax

Likewise, "<value expression> <comp op> ALL <table expression>" takes each row, *s*, of <table expression> and tests <value expression> <comp op> *s* with ANDs between the expanded expressions. We generalized the logic what are called quantifiers.

```
((<value expression> <comp op> s1)
AND (<value expression> <comp op> s2)
...
AND (<value expression> <comp op> sn))
```

When you get a single FALSE result, the whole predicate is FALSE. As long as <table expression> has cardinality greater than zero and all non-NULL values, you will get a result of TRUE or FALSE.

That sounds reasonable so far. Now let `Empty_Set` be an empty table (no rows, cardinality zero) and `NullTable` be a table with only NULLs in its rows and a cardinality greater than zero. The rules for SQL say that "<value expression> <comp op> ALL `NullTable`" always returns UNKNOWN, and likewise "<value expression> <comp op> ANY `NullTable`" always returns UNKNOWN. This makes sense, because every row comparison test in the expansion would return UNKNOWN, so the series of OR and AND operators would behave in the usual way.

However, "<value expression> <comp op> ALL `Empty_Set`" always returns TRUE and "<value expression> <comp op> ANY `Empty_Set`" always returns FALSE. Most people have no trouble seeing why the ANY predicate works that way; you cannot find a match, so the result is FALSE. But most people have lots of trouble seeing why the ALL predicate is TRUE. This convention is called existential import, and it was a big debate at the start of modern logic. It boiled down to deciding if the statement "All *x* are *y*" implies that "some *x* exists" by definition. The modern convention is that it does not. Lewis Carroll went with existential import, John Venn and George Boole did not. And neither did modern logic; let me explain.

If I were to walk into a bar and announce that I can beat any pink elephant in the bar, that would be a true statement. The fact that there are no pink elephants in the bar merely shows that the problem is reduced to the minimum case. If this seems unnatural, then convert the ALL and ANY predicates into EXISTS predicates and look at the way that this rule preserves the formal mathematical properties that:

$$1) (\forall x)P(x) = \neg(\exists x) \neg P(x)$$

$$2) (\exists x)P(x) = \neg(\forall x)\neg P(x)$$

If you have not had a mathematical logic course or just forgot everything, the $(\forall x)$ means “for all x ..” and $(\exists x)$ means “there exists one or x, such that ..” and these rules can be used by an optimizer (with allowances for NULL and three valued logic).

The ALL Predicate and Extrema Functions

It is counter-intuitive at first that these two predicates are not the same in SQL:

```
x >= (SELECT MAX(y) FROM Table1)
x >= ALL (SELECT y FROM Table1)
```

But you have to remember the rules for the extrema functions – they drop out all the NULLs before returning the greatest or least values. The ALL predicate does not drop NULLs, so you can get them in the results.

However, if you know that there are no NULLs in a column or are willing to drop the NULLs yourself, then you can use the ALL predicate to construct single queries to do work that would otherwise be done by two queries. For example, given the table of products and store managers, to find which manager handles the largest number of products, you would first construct a grouped VIEW or CTE and group it again. Assume we have a table of the managers responsible for each product we stock:

Now, the query:

```
WITH TotalProducts (manager_name, product_tally)
AS
  (SELECT manager_name, COUNT(*)
   FROM Product_Managers
   GROUP BY manager_name)

SELECT manager_name
   FROM TotalProducts
  WHERE product_tally
         = (SELECT MAX(product_tally)
           FROM TotalProducts);
```

But Alex Dorfman found a single query solution instead:

```
SELECT manager_name, COUNT(*) AS product_tally
  FROM Product_Managers
 GROUP BY manager_name
HAVING COUNT(*) + 1
       > ALL (SELECT DISTINCT COUNT(*)
             FROM Product_Managers
             GROUP BY manager_name);
```

The use of the SELECT DISTINCT in the subquery is to guarantee that we do not get duplicate rows when two managers handle the same number of products. You can also add a ".. WHERE dept IS NOT NULL" clause to the subquery to get the effect of a true MAX() aggregate function.

EXISTS()

The EXISTS predicate is very natural. It is a test for a non-empty table. If there are any rows in its subquery, it is TRUE; otherwise, it is FALSE; this predicate does not give an UNKNOWN result. The syntax is simply EXISTS <table subquery>. It is worth mentioning that a <table subquery> is always inside parentheses to avoid problems in the grammar during parsing.

Some early SQL implementations would work better with "EXISTS(SELECT <column name> ..)", " EXISTS(SELECT <constant> ..)" or "EXISTS(SELECT * ..)" versions of the predicate. Today, there is no difference in the three forms in the major products, so the "EXISTS(SELECT * ..)" is now the preferred form since it shows that we are *working at the table and row level*, without regard to columns.

Indexes are very useful for EXISTS() predicates because they can be searched while the base table is left completely alone. For example, we want to find all the employees who were born on the same day as any famous person. The query could be

```
SELECT P.emp_name AS famous_person_birth_date_guy
  FROM Personnel AS P
 WHERE EXISTS
    (SELECT *
     FROM Celebrities AS C
     WHERE P.birth_date = C.birth_date);
```

If the table `Celebrities` has an index on its `birth_date` column, the optimizer will get the current employee's `birth_date`, `P.birth_date`, and look up that value in the index. If the value is in the index, the predicate is `TRUE` and we do not need to look at the `Celebrities` table at all. If it is not in the index, the predicate is `FALSE` and there is still no need to look at the `Celebrities` table. This should be fast, since indexes are smaller than their tables and are structured for very fast searching.

However, if `Celebrities` has no index on its `birth_date` column, the query may have to look at every row to see if there is a `birth_date` that matches the current employee's `birth_date`. There are some tricks that a good optimizer can use to speed things up in this situation.

A `NULL` might not be a value, but it does exist in SQL. This is often a problem for a new SQL programmer who is having trouble with `NULLs` and how they behave.

Think of them as being like a brown paper bag -- you know that something is inside because you lifted it up and felt a weight, but you do not know exactly what that something is. If you felt an empty bag, you know to stop looking. For example, we want to find all the `Personnel` who were not born on the same day as a famous person. This can be answered with the negation of the original query, like this:

```
SELECT P.emp_name AS famous_birth_date_person
FROM Personnel AS P
WHERE NOT EXISTS
  (SELECT *
   FROM Celebrities AS C
   WHERE P.birth_date = C.birth_date);
```

But assume that among the `Celebrities`, we have a movie star who will not admit her age, shown in the row ('Gloria Glamour', `NULL`). A new SQL programmer might expect that Ms. Glamour would not match to anyone, since we do not know her `birth_date` yet. Actually, she will match to everyone, since there is a chance that they may match when some tabloid newspaper finally gets a copy of her birth certificate. But work out the subquery in the usual way to convince yourself:

Another problem with `NULLs` is found when you attempt to convert `IN()` predicates to `EXISTS` predicates. Using our example of matching our `Personnel` to famous people, the query can be rewritten as:

```
SELECT P.emp_name AS famous_birth_date_person
FROM Personnel AS P
WHERE P.birth_date
      NOT IN
      (SELECT C.birth_date
       FROM Celebrities AS C);
```

However, consider a more complex version of the same query, where the celebrity has to have been born in New York City. The IN() predicate would be

```
SELECT P.emp_name -- born on a day without a famous New Yorker!
FROM Personnel AS P
WHERE P.birth_date
      NOT IN
      (SELECT C.birth_date
       FROM Celebrities AS C
       WHERE C.birth_city_name = 'New York');
```

and you would think that the EXISTS version would be:

```
SELECT P.emp_name, -- born on a day without a famous New Yorker!
FROM Personnel AS P
WHERE NOT EXISTS
      (SELECT *
       FROM Celebrities AS C
       WHERE C.birth_city_name = 'New York'
       AND C.birth_date = P.birth_date);
```

Assume that Gloria Glamour is our only New Yorker and we still do not know her birth_date. The subquery will be empty for every employee in the NOT EXISTS predicate version, because her NULL birth_date will not test equal to the known employee birthdays.

That means that the NOT EXISTS predicate will return TRUE and we will get every employee to match to Ms. Glamour. But now look at the IN() predicate version, which will have a single NULL in the subquery result. This predicate will be equivalent to (Personnel.birth_date = NULL), which is always UNKNOWN, and we will get no Personnel back.

Likewise, you cannot, in general, transform the quantified comparison predicates into EXISTS predicates, because of the possibility of NULL values. Remember that "x <> ALL <subquery>" is shorthand for "x NOT IN() <subquery>" and "x = ANY <subquery>" is shorthand for "x IN <subquery>" and it will not surprise you.

In general, the EXISTS predicates will run faster than the IN() predicates. The problem is in deciding whether to build the query or the subquery first; the optimal approach depends on the size and distribution of values in each, and that cannot usually be known until run-time.

The OVERLAPS() Predicate

The OVERLAPS() predicate for finding out if two temporal periods overlap each. Let's call the periods (S1, T1) and (S2, T2). The S is for the starting timestamp of an event, T is for the termination timestamp of the event. The result of the <OVERLAPS predicate> is formally defined as the result of the following expression:

```
(S1 > S2 AND NOT (S1 >= T2 AND T1 >= T2))  
OR (S2 > S1 AND NOT (S2 >= T1 AND T2 >= T1))  
OR (S1 = S2 AND (T1 <> T2 OR T1 = T2))
```

The rules for the OVERLAPS() predicate sound like they should be intuitive, but they are not. The principles that we wanted in the Standard were:

1. A time period includes its starting point, but does not include its end point. We have already discussed this model and its closure properties.
2. If the time periods are not "instantaneous", they overlap when they share a common time period.
3. If the first term of the predicate is an INTERVAL and the second term is an instantaneous event (a <datetime> data type), they overlap when the second term is in the time period (but is not the end point of the time period). That follows the half-open model.
4. If the first and second terms are both instantaneous events, they overlap only when they are equal.
5. If the starting time is NULL and the finishing time is a <datetime> value, the finishing time becomes the starting time and we have an event. If the starting time is NULL and the finishing time is an INTERVAL value, then both the finishing and starting times are NULL.

Please consider how your intuition reacts to these results, when the granularity is at the YEAR-MONTH-DAY level. Remember that a day begins at 00:00:00 Hrs.

```
(today, today) OVERLAPS (today, today) = TRUE
(today, tomorrow) OVERLAPS (today, today) = TRUE
(today, tomorrow) OVERLAPS (tomorrow, tomorrow) = FALSE
(yesterday, today) OVERLAPS (today, tomorrow) = FALSE
```

The full definition is more complex, but the point is that SQL can hide the complexity with a single syntactic construct.

The CASE Expression

The CASE expression is probably the most useful addition made in the SQL-92 Standard. ANSI stole the idea and the syntax from the ADA programming language and the selection operators from recursive function theory. Here is the BNF for a <case specification>:

```
<case specification> ::= <simple case>, <searched case>
<simple case> ::=
CASE <case operand>
<simple when clause>...
[<else clause>]
END

<searched case> ::=
CASE
<searched when clause>...
[<else clause>]
END

<simple when clause> ::= WHEN <when operand> THEN <result>
<searched when clause> ::= WHEN <search condition> THEN
<result>
<else clause> ::= ELSE <result>
<case operand> ::= <value expression>
<when operand> ::= <value expression>
<result> ::= <result expression>, NULL
<result expression> ::= <value expression>
```

It can be used to move old procedural code into declarative code that can be optimized. If you can use the <simple case>, do it. It will optimize easily. While it is logically equivalent to a searched case with only equality tests on the <case operand>, nobody implements it that way. It is much faster to use hashing and other techniques for longer CASE expressions.

The next trick is to remember that the WHEN clauses are tested in left-to-right order. If you arrange them so that the most likely ones appears first, you will not waste time with the rare events. In medicine, this principle is expressed as “Look for a horse, not a zebra” when doing diagnostics.

Most optimizers will also factor out redundant tests by considering the order of execution. That is easier to see with an example:

```
CASE WHEN x = 1 THEN 'alpha'  
      WHEN x = 1 AND y = 5 THEN 'beta'  
      ELSE 'gama' END
```

if ($x = 1$) then the first WHEN clauses will execute; there is no need to re-test that condition in the following WHEN clauses. If it worked, then we never got any further than the first THEN clause.

However, most optimizers are weak on nested CASE expressions. This is funny because we spent a lot of time with nested if-then-else constructs in procedural languages. It is probable better to flatten out nested CASE expressions to a single level.

Do Not Retrieve More Columns Than You Need

Blindly using “SELECT *” in a query is a waste of resources, since you will probably not need everything in the result set. You are also locking resources from other users; this is especially important in columnar SQL databases.

But you can get extraneous columns other ways. When you do a “cut & paste” to get some code re-use you can get more than you need. Remember to look at actual text. When you invoke a VIEW with more columns than needed, you are creating the same problem. If the VIEWS are nested very deeply, you probably cannot see the extraneous columns. You should un-nest the poorly performing query and see everything. Look for a few over-sized columns that are not used (did you really need a full product description or just the UPC code?).

Symptoms of this problem is a WHERE clause with extra conditions and a FROM clause with extra outer joins.

DELETE and UPDATE in Batches

Deleting or updating large amounts of data from a huge table scan can eat resources. The problem is that both of these statements run as a single transaction, and if you need to kill them or if something happens to the system while they are working, the system has to roll back the entire transaction. This consumes lots of time, and blocks other transactions.

The solution is to do deletes or updates in smaller batches. If the transaction gets killed, the ROLLBACK is relatively small. And while each small batch is committing to disk, other transactions can “sneak in” and do some work, so concurrency is greatly enhanced. If you are archiving, you might want to take a little longer than a production jobs.

Check your SQL; you may find that you have CHECKPOINT or SAVEPOINT (the name varies) features. These are like “mini-commits” which let a localized ROLLBACK restore the database back to the SAVEPOINT, or COMMIT the work done so far. This might be a variable option in your situation.

Use Stored Procedures

Stored procedures have a number of advantages. The parameters are parsed and verified easily by the SQL engine. Less data is pushed across the network. The stored procedure call will always be much shorter. Also, stored procedures are easier to trace in profile tools. A stored procedure is an actual object in your database. That means it is much easier to get performance statistics on a stored procedure than on an ad hoc query and, in turn, find performance issues and draw out anomalies.

In addition, stored procedures parameterize more consistently. This means you are more likely to reuse your execution plans and even deal with caching issues, which can be difficult to pin down with ad hoc queries. Stored procedures also make it much easier to deal with edge cases and even add auditing or change-locking behavior. A stored procedure can handle many tasks that trouble ad hoc queries.

Stored procedures are also not subject to SQL injections attacks. Take a look at this cartoon (<https://xkcd.com/327/>) to get a feeling as to how SQL injection works. The attacker passes a string of code as an argument to a block of dynamic SQL. This string contains malicious SQL statements that are executed when the dynamic SQL is prepared.

Finally, when you find a better way to do something, you just replace the procedure body. If you can do the task with a subset of the old parameter list, then life is good. If you need to add a new parameter, then you have to do some work.

GUIDs are Not What You Think

A UUID (universal unique identifier) or a GUID (globally unique identifier) is a 16-byte randomly generated number. It is not an attribute of anything in the database, so by definition, it cannot be a key. A key is a subset of attribute of the an entity. When you use this, you are faking a physical pointer and not doing RDBMS.

Ordering your table's data on this column will cause your table to fragment much faster than using a steadily increasing value like a DATE or SEQUENCE.

The "G" in GUID means "Global" and should never be used for local data. It was created to locate external objects. Instead look for a real identifier that has validations and verification rules.

Do Not Use Count() if You Need Existence

This is a common error. Using a simple count to test for existence is insane.

```
IF (SELECT COUNT(*) FROM Foobars) > 0
THEN
<database actions>
END IF;
```

it is completely unnecessary. If you want to check for existence, then do this:

```
IF EXISTS (SELECT * FROM Foobars)
THEN
<database actions>
END IF;
```

Do not count everything in the table. Just get back the first row you find. SQL is smart enough to use EXISTS() properly, and the second block of code returns much faster. The larger the table, the bigger difference this will make.

Do Not Do Negative Searches

Take the simple query

```
SELECT *
FROM Customers
WHERE region_nbr <> 3;
```

You cannot use an index with this query because it is a negative search that has to be compared row by row with a table scan. If you need to do something like this, you may find it performs much better if you rewrite the query to use the index.

Use CTEs

In linguistics and mathematics, the ability to replace a simple element with a complex element of the same type is called orthogonality. A noun in a sentence can be replaced with a noun phrase, in a formula, an integer can be replaced by an integer function, and so forth.

If you had a good grammar class in middle school, you might have seen this sentence "The rat the cat the dog chased killed ate the malt"; it is perfectly correct. It is a sentence with multiple center embeddings. We usually have no problem putting one clause inside another in English. We can take "the rat ate the malt" and stick in more information to make "the rat the cat killed ate the malt." But the more clauses we add in, the harder it gets to understand the sentence. In this case, the rat ate the malt. After that it was killed by a cat. That cat had been chased by a dog. The grammar of the sentence is fine. The style and readability are awful.

With a CTE, you have moved a derived table expression from the interior of SELECT statement to the front of it and put a name on it. The CTEs are created in left to right order, which means that a CTE can only reference previous declared CTEs.

However, once a CTE exists, it can be referenced by name in multiple places in the main statement. In theory a good optimizer would find multiple occurrences of the same code, factor it out and decide if the code should be treated like a macro (i.e. copied into the program text, then compiled everywhere it is used) or if it should be

materialized as a working table one time. The theory is nice, but deep nesting can fool an optimizer into taking the safe route and using the macro approach

The worst such construct I have seen is 16 pages of CTEs leading to a single SELECT * statement on the last page. Think about what nesting them as derived tables would look like, however! That would have been much worse. If the CTE gets a good name, you can quickly see if it should be made into a VIEW and shared with other queries.

“Double-Dipping” and the Windows Clause.

This means avoid visiting the same table more than once. A horrible way to do this is with temporary tables built from the same base table, then joined together. If your SQL engine's optimizer can decide to materialize VIEWS and share the results as a working table, then several sessions can benefit. If you have a less sophisticated optimizer, you can use temp tables and do this work by hand. a set.

The Window clause is the best one of the best ways to avoid this and the performance improvements can be huge. The subclauses are similar to features of a subquery. The PARTITION BY is like a local GROUP BY that splits up the table into groupings. The ORDER BY imposes a sorted order. Finally, the window frame (ROW or RANGE) is like a local WHERE clause.

This is a complex construct and you need to spend time with it. To get an idea of the power, let's do a common, simple running total. Many report writers have it built in, but in SQL, before the windows clause, we had to use self-joins for this. Assume we have a table of Sales

```
CREATE TABLE Sales
(sales_invoice_nbr CHAR(5) NOT NULL PRIMARY KEY,
 emp_id CHAR(10) NOT NULL
REFERENCES Personnel(emp_id),
 sales_date DATE NOT NULL,
 sales_amt DECIMAL(10,2) NOT NULL);
```

To do a simple running total, ordered by the date of the sales

```
SELECT S1.sales_date, SUM(S2.sales_amt) AS sales_amt_runtot
FROM Sales AS S1, Sales AS S2
WHERE S2.sales_date <= S1.sales_date
GROUP BY S1.sales_date
ORDER BY S1.sales_date; -- not really needed
```

The table appears as the row with the current sales, S1, and a second copy, S2, of the table appears as all *prior* sales, we a self-join, apply the WHERE clause and then group the rows in the result set. Each row in S1 invokes a query with a join to S2. Over and over and over.

But with a windows clause, we can write;

```
SELECT sales_date, emp_id,
       SUM(sales_amt)
         OVER (ORDER BY sales_date
              ROWS BETWEEN UNBOUNDED PRECEDING AND CURRENT ROW)
       AS sales_amt_runtot
FROM Sales;
```

This uses only one copy of the Sales table. The extra keywords tell the optimizer that it can build an ordered list, total it and then for each row add the current sales_amt. Mere self-joins cannot tell the optimizer about the *intent* of the query. If you want to get the running totals by salesman, just a partition clause:

```
SELECT sales_date, emp_id,
       SUM(sales_amt)
         OVER (PARTITION BY emp_id
              ORDER BY sales_date
              ROWS BETWEEN UNBOUNDED PRECEDING AND CURRENT ROW)
       AS sales_amt_runtot
FROM Sales;
```

The full syntax of the window clause is elaborate and not fully implemented in many SQLs. The basic types of windowed functions are and aggregate functions, rank functions and position functions.

If the window starts at UNBOUNDED PRECEDING, then the lower bound is always the first row of the partition; or you can give an actual count of rows to use for the frame. CURRENT ROW explains itself. The <window frame preceding> is an actual count of preceding rows.

UNBOUNDED FOLLOWING or an actual count of following rows are the same ideas on the other side of the current row. There are also <window frame exclusion> clauses that remove the current :row, a group and tires.

The aggregate functions are the usual ones from earlier SQL standards – SUM(), MIN(), MAX(), AVG() and COUNT(). We already saw an example of a summation and the other aggregates work the same way. Obviously, the ORDER BY clause has no effect on MIN(), MAX() and COUNT()

The rank functions are ROW_NUMBER(), RANK(), DENSE_RANK(), PERCENT_RANK(), and CUME_DIST(). They assign an ordering number within the scope of the window. ROW_NUMBER() is the most important and the most obvious; it returns a number from 1 to (n) for all of that rows in the window. Again there are a lot of programming idioms with multiple calls to this function.

The position functions are not yet part of the standards at this time, but they are very common. There is LEAD(<column>[<n>]) and LAG(<column>[<n>]) which return the value in a column that is (n) rows preceding or following the current row. There are also FIRST and LAST functions that work with the rows at the edges of the window.

It is worth mentioning that the rules for an ORDER BY subclause have changed to be more general than they were in earlier SQL Standards. A sort can now be a <value expression> and is not limited to a simple column in the select list. However, it is still a good idea to use only column names so that you can see the sorting order in the result set. You put NULL either first or last as well as ASC and DESC options.

Conclusions

Just as took Mr. Sedaris some time to learn “talk pretty one day” when he went to live in France, it will take some time for new sql developers to “talk pretty” and to get out of the mindset of their prior language. If I had to summarize all of the tips in this article, it would be to write SQL that looks like SQL, just as you would try to speak French that sounds like real French instead of the “Fran-lish” that comes from reading a phrase book with an American accent.

Learn each SQL statement and all of the tricks inside it, even if you have no immediate need.

Practice coding SQL. Do not expect to have an epiphany overnight, but one day you will talk pretty in SQL.

About The Author

Mr. Joe Celko serves as Member of Technical Advisory Board of Cogito, Inc. Mr. Celko joined the ANSI X3H2 Database Standards Committee in 1987 and helped write the ANSI/ISO SQL-89 and SQL-92 standards. He is one of the top SQL experts in the world, writing over 700 articles primarily on SQL and database topics in the computer trade and academic press. The author of six books on databases and SQL, Mr. Celko also contributes his time as a speaker and instructor at universities, trade conferences and local user groups.

Download a Free Trial at www.embarcadero.com

Corporate Headquarters | Embarcadero Technologies | 275 Battery Street, Suite 1000 | San Francisco, CA 94111 | www.embarcadero.com | sales@embarcadero.com

© 2015 Embarcadero Technologies, Inc. Embarcadero, the Embarcadero Technologies logos, and all other Embarcadero Technologies product or service names are trademarks or registered trademarks of Embarcadero Technologies, Inc.
All other trademarks are property of their respective owners 100815